Edge-Intelligent Industrial IoT for Occupational-Safety Monitoring and Automated Compliance Reporting

Rishwanth Balakrishnan
B.E., Dept. of Electronics and Communication Engineering
Sathyabama Institute of Science and Technology
Chennai, India

Rakesh M
B.E., Dept. of Electronics and Communication Engineering
Sathyabama Institute of Science and Technology
Chennai, India

Dr. S. Poornapushpakala
M.E., Ph.D., Professor (Guide), Dept. of Electronics and Communication Engineering
Sathyabama Institute of Science and Technology
Chennai, India

***Abstract*—Occupational hazards such as toxic and combustible gases, airborne dust, excessive noise and unsafe thermal conditions continue to cause a large share of industrial injuries and long-term occupational disease, yet the dominant safety practice in most facilities remains discrete and manual. This survey reviews the body of work that moves occupational-safety monitoring from periodic human inspection toward continuous, networked, and increasingly autonomous sensing. We organise the literature around a four-stage evolution from manual practice through sensor-based standalone detectors and cloud-centric Internet-of-Things (IoT) platforms to edge-first architectures, and examine the enabling components at each stage: low-cost gas and environmental sensing, microcontroller platforms, lightweight publish/subscribe communication, and on-node intelligence. Particular attention is given to two capabilities that distinguish recent edge-intelligent designs from earlier threshold detectors: the computation of cumulative-dose metrics such as the Time-Weighted Average (TWA), and the generation of tamper-evident, auditor-ready compliance records. We compare the paradigms along the axes of latency, availability, scalability and evidentiary value, summarise representative sensing channels and their statutory limits, and identify open challenges in sensor calibration, energy autonomy, security and standardised reporting. The survey is intended as a structured entry point for researchers and practitioners designing plant-scale occupational-safety systems from accessible, low-cost hardware.**

***Index Terms*—Industrial Internet of Things, occupational safety, edge computing, gas sensing, MQTT, time-weighted average, compliance reporting, ESP32.**

I. INTRODUCTION

Industrial environments such as foundries, chemical plants, boiler rooms, paint shops, food-processing units and warehouses are among the most hazardous workplaces in the modern economy. Workers are routinely exposed to toxic and flammable gases, airborne dust, excessive noise and thermally unsafe conditions. A leak of carbon monoxide or liquefied petroleum gas can incapacitate within minutes, while chronic low-level exposure produces long-term disease that is far harder to attribute and prosecute. Despite decades of regulation, a large fraction of industrial units, particularly in developing regions, still depend on periodic manual inspection using handheld meters, paper logbooks and the vigilance of individual supervisors.

The central weakness of this practice is that it is discrete rather than continuous: a concentration measured at one instant says nothing about a leak that develops minutes later, and because many hazards are invisible, odourless or intermittent, the interval between manual checks is precisely the interval during which workers are unprotected. A second weakness is the absence of trustworthy records; paper logs are easily lost, back-filled or written after the fact, and therefore cannot demonstrate the true exposure profile of a shift. A third is scale: a large facility contains many distinct hazard zones, yet a supervisor can be in only one place at a time.

The Internet of Things (IoT) offers a structural response to these limitations by embedding networked sensing and computation directly into the workplace [1], [2]. A dense mesh of low-cost nodes can watch every zone simultaneously and without fatigue. However, the manner in which intelligence is distributed across such a mesh, whether cloud-centric or edge-first, has a decisive effect on latency, resilience and evidentiary value [4], [5]. This survey reviews the literature underpinning that design space, with an emphasis on edge-intelligent architectures that retain deterministic local protection while still delivering scalable oversight and audit-quality records.

The remainder of the paper is organised as follows. Section II characterises the hazards and the monitoring requirements they impose. Section III presents a four-stage taxonomy of occupational-safety monitoring. Sections IV through VIII survey the enabling technologies: sensing, communication, edge intelligence, exposure analytics and compliance reporting. Section IX discusses open challenges and future directions, and Section X concludes.

II. OCCUPATIONAL HAZARDS AND MONITORING REQUIREMENTS

The hazards found on an industrial floor fall into a small number of broad categories, each demanding a dedicated sensing channel. Toxic gases, chief among them carbon monoxide (CO), are colourless and odourless and act on a time scale of minutes. Combustible gases and vapours, such as liquefied petroleum gas (LPG) and solvent fumes, pose an explosion risk once their concentration approaches the lower explosive limit (LEL). Airborne particulates and poor air quality contribute to chronic respiratory disease, while physical stressors such as excessive noise, heat and inadequate lighting degrade both health and situational awareness over time [12].

These hazards impose three monitoring requirements that jointly define the design problem. First, timeliness: acute hazards must be detected and responded to within seconds, which bounds acceptable end-to-end latency. Second, cumulative accounting: chronic hazards are governed not by instantaneous peaks but by dose accumulated over a shift, so regulatory frameworks specify a Time-Weighted Average (TWA) alongside instantaneous limits [14], [15]. Third, evidentiary integrity: because safety compliance is a legal obligation, a monitoring system must produce records that can survive audit and incident investigation. A system that satisfies only the first requirement is a mere alarm; satisfying all three is what separates a modern occupational-safety platform from a simple detector.

III. EVOLUTION OF OCCUPATIONAL-SAFETY MONITORING

The literature can be read as a progression through four overlapping stages, each retaining the strengths of its predecessor while addressing a residual weakness. Table I summarises the comparison.

TABLE I
Comparison of Occupational-Safety Monitoring Paradigms

Attribute	**Manual practice**	**Standalone detector**	**Cloud-centric IoT**	**Edge-first IoT**
Sensing continuity	Discrete (rounds)	Continuous, local	Continuous, networked	Continuous, networked
Decision latency	Minutes to hours	Milliseconds (local)	Network-dependent	Milliseconds (local)
Availability on network loss	N/A	Full (isolated)	Lost	Full (local)
Plant-wide visibility	None	None	Yes	Yes
Cumulative dose (TWA)	Rare/manual	Seldom	Possible (server)	On-node
Auditable records	Weak (paper)	Weak	Moderate	Strong (tamper-evident)
Cost per zone	Low (labour-heavy)	Low	Moderate	Low–moderate

Manual practice relies on handheld detectors and scheduled rounds. It is inexpensive and simple but discrete, labour-intensive and record-poor. Sensor-based standalone systems embed fixed detectors, frequently built around low-cost metal-oxide gas sensors of the MQ family and a microcontroller, which provide continuous sensing and autonomous local response such as an alarm or a shutoff [13]. They remove the gaps between rounds but remain isolated, single-purpose

and unable to give a plant-wide picture. Cloud-centric IoT systems connect many nodes wirelessly to a remote server that stores data, visualises it on a dashboard and issues alerts [3]. This introduces networking, remote visibility and historical storage and represents the current state of the art in many commercial products; its principal weakness is architectural, in that detection and alerting depend on an uninterrupted network path to the server [6], [7].

Edge-first architectures relocate time-critical intelligence to the node itself [4], [8]. Threshold comparison, TWA accumulation, local alarm and ventilation actuation are performed on-device, so protection does not depend on the cloud; the network is reserved for oversight, aggregation and analytics. This division of labour, in which fast autonomous protection sits at the edge and scalable oversight at the centre, is the defining principle of recent designs and the reason such systems degrade gracefully: if the broker, dashboard or wider network fails, each node continues to sense, decide and protect locally.

IV. SENSING TECHNOLOGIES FOR INDUSTRIAL HAZARD DETECTION

The accessibility of edge-intelligent safety systems rests largely on the availability of inexpensive sensing elements. Surveys of low-cost gas sensing distinguish five dominant transduction principles: electrochemical, catalytic, semiconductor (metal-oxide), non-dispersive infrared (NDIR) and photo-ionisation. Each is sensitive to a particular class of gases and none is universal [12]. In the low-cost, multi-parameter designs that dominate the recent edge-IoT literature, semiconductor sensors of the MQ family are ubiquitous: the MQ-7 for carbon monoxide, the MQ-6/MQ-2 for LPG and combustible vapours, and the MQ-135 for general air quality [13]. Systematic reviews of IoT gas-detection work report that these sensors, paired with Xtensa-architecture microcontrollers such as the ESP32/ESP8266 and, for heavier processing, ARM Cortex-A single-board computers, constitute the most common hardware pattern in the field. Table II lists representative channels and their governing limits.

TABLE II
Representative Low-Cost Sensing Channels and Target Hazards

Channel	Typical device	Hazard / metric
Carbon monoxide	MQ-7 (metal-oxide)	Acute toxicity; TWA in ppm
LPG / combustible	MQ-6 / MQ-2	Explosion risk; % of LEL
Air quality	MQ-135	Chronic exposure; ppm
Particulates	Optical dust sensor	Respiratory disease
Noise	Electret / MEMS mic	Hearing loss; dB(A)
Temp./humidity	DHT / SHT series	Heat stress; comfort
Illumination	Photodiode / LDR	Visibility; ergonomics

Metal-oxide sensors are attractive for their low cost, fast response and simple analog interface, but they exhibit cross-sensitivity, baseline drift and temperature/humidity dependence. Consequently the literature stresses calibration, compensation and, increasingly, learned correction as prerequisites for trustworthy readings, a point returned to in Section IX. Complementary channels for dust, noise, temperature, humidity and illumination extend a node from a single-gas detector into a multi-hazard occupational-safety instrument.

V. COMMUNICATION PROTOCOLS AND NETWORK ARCHITECTURES

Once sensing is distributed across many nodes, the choice of communication protocol governs scalability and energy cost. Comparative studies of the application-layer IoT protocols MQTT, CoAP, AMQP and HTTP consistently favour the publish/subscribe model for many-to-one telemetry because of its low header overhead and decoupling of producers from consumers [2], [11]. MQTT, standardised by OASIS [10], has become the de facto choice for plant-scale sensor telemetry: nodes publish readings to a hierarchical topic namespace, a broker fans them out to

dashboards and analytics subscribers, and Quality-of-Service levels allow a design to trade bandwidth for delivery guarantees.

At the network layer, early industrial monitoring systems favoured ZigBee/IEEE 802.15.4 mesh topologies for their low power and self-healing routing [12], whereas contemporary low-cost nodes frequently adopt Wi-Fi for its ubiquity and direct IP connectivity. A recurring architectural theme is the separation of a fast local control path (sensor → node logic → actuator) from a slower telemetry path (node → broker → dashboard), so that communication latency or loss never sits on the critical safety path [6], [7].

VI. EDGE COMPUTING AND ON-NODE INTELLIGENCE

Edge computing, which performs computation near the data source rather than in a distant cloud, directly addresses the latency and availability weaknesses of cloud-centric designs [4], [5]. In the occupational-safety context the intelligence relocated to the node is modest but decisive: real-time threshold evaluation against statutory limits, hysteresis to suppress nuisance alarms, running TWA accumulation, and the sequencing of local responses such as status indication, audible alarm and automatic ventilation. Because these decisions are taken on-device, the system provides deterministic protection that is independent of network state, and a mesh of such nodes exhibits graceful degradation rather than a single point of failure [9].

A growing strand of work pushes further, embedding lightweight machine learning on the microcontroller itself (the "TinyML" paradigm) to classify abnormal environmental patterns, compensate sensor drift, or anticipate hazardous trends before fixed thresholds are crossed [8], [17]. On-node inference keeps raw data local, reduces uplink traffic, and preserves function during outages, at the cost of tight memory and energy budgets that shape model selection and quantisation.

VII. EXPOSURE ANALYTICS: TIME-WEIGHTED AVERAGE AND PREDICTIVE METHODS

A feature that distinguishes modern safety platforms from simple detectors is their treatment of cumulative exposure. Regulatory limits such as the ACGIH Threshold Limit Values and OSHA Permissible Exposure Limits are defined in part as time-weighted averages over a reference period, classically an eight-hour shift [14], [15]. The TWA over a period T for a concentration signal $c(t)$ is

$$\mathrm{TWA} = (1/T) \int_0^T c(t)\, \mathrm{d}t \approx (1/T) \sum_i c_i\, \Delta t_i \tag{1}$$

where the discrete form is what a node actually evaluates by integrating periodic samples c_i over their sampling intervals Δt_i. Computing this quantity on-node lets the system model a worker's dose over a shift and, crucially, distinguish a brief acute spike, which demands evacuation, from slow chronic accumulation, which demands intervention but not panic. Computing TWA by hand from sparse manual samples is effectively impossible, which is why the metric is a natural beneficiary of continuous edge sensing. Beyond fixed integration, predictive approaches apply trend analysis and learned models to forecast threshold breaches and issue early warnings [8].

VIII. COMPLIANCE REPORTING AND DATA INTEGRITY

The final capability surveyed is automated compliance reporting. Because paper records are weak evidence, recent systems generate machine-written, timestamped logs of conditions and actions, and increasingly seek to make those logs tamper-evident so that they can support audit and incident investigation. Techniques discussed in the broader IoT-security literature, such as hash-chaining of log entries, signed records and append-only stores, carry directly into this setting [2]. Coupled with multi-channel, escalating alerts to safety personnel, auto-generated compliance reports close the loop between sensing and organisational accountability, converting raw sensor values into both immediate protective action and durable evidence. This evidentiary dimension, largely absent from standalone detectors, is where edge-first architectures offer their clearest advantage over prior paradigms (Table I).

IX. OPEN CHALLENGES AND FUTURE DIRECTIONS

Several challenges recur across the surveyed literature. *Sensor accuracy and calibration:* low-cost metal-oxide sensors drift and cross-respond, so field calibration, environmental compensation and learned correction remain active problems that bound the trust one can place in cheap nodes [12]. *Energy autonomy:* continuous sensing, local computation and wireless telemetry must fit within battery or energy-harvesting budgets, particularly for outdoor or mobile nodes. *Security and integrity:* as safety records acquire legal weight, the confidentiality, authenticity and non-repudiation of both telemetry and logs become essential, drawing on the edge/fog security literature [6]. *Standardisation and interoperability:* the field is fragmented across ad-hoc topic schemes, dashboards and report formats, and would benefit from common data models for exposure and compliance. *Validation:* finally, most low-cost designs are demonstrated on prototypes; rigorous, longitudinal validation against reference instruments in real industrial settings is comparatively rare and is arguably the most important direction for future work.

X. CONCLUSION

This survey has traced occupational-safety monitoring from discrete manual practice to continuous, edge-intelligent Industrial IoT. Organising the literature as a four-stage evolution clarifies why edge-first architectures have become the reference design: they preserve the deterministic, locally autonomous protection of standalone detectors while adding the plant-wide visibility of cloud IoT, and they uniquely support on-node cumulative-dose analytics and tamper-evident compliance reporting. The enabling ingredients (low-cost multi-parameter sensing, capable microcontrollers, lightweight publish/subscribe messaging and on-device intelligence) are now individually mature; the open problems lie in calibration, energy autonomy, security, standardised reporting and field-scale validation. Addressing them would allow robust, plant-scale safety platforms to be built from accessible hardware, extending proactive protection to precisely the smaller, less-regulated facilities where the human cost of occupational hazards is highest.

REFERENCES

[1] J. Gubbi, R. Buyya, S. Marusic, and M. Palaniswami, “Internet of Things (IoT): A vision, architectural elements, and future directions,” *Future Gener. Comput. Syst.*, vol. 29, no. 7, pp. 1645–1660, 2013. https://doi.org/10.1016/j.future.2013.01.010

[2] A. Al-Fuqaha, M. Guizani, M. Mohammadi, M. Aledhari, and M. Ayyash, “Internet of Things: A survey on enabling technologies, protocols, and applications,” *IEEE Commun. Surveys Tuts.*, vol. 17, no. 4, pp. 2347–2376, 2015. https://doi.org/10.1109/COMST.2015.2444095

[3] A. Zanella, N. Bui, A. Castellani, L. Vangelista, and M. Zorzi, “Internet of Things for smart cities,” *IEEE Internet Things J.*, vol. 1, no. 1, pp. 22–32, 2014. https://doi.org/10.1109/JIOT.2014.2306328

[4] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, “Edge computing: Vision and challenges,” *IEEE Internet Things J.*, vol. 3, no. 5, pp. 637–646, 2016. https://doi.org/10.1109/JIOT.2016.2579198

[5] M. Satyanarayanan, “The emergence of edge computing,” *Computer*, vol. 50, no. 1, pp. 30–39, 2017. https://doi.org/10.1109/MC.2017.9

[6] F. Bonomi, R. Milito, J. Zhu, and S. Addepalli, “Fog computing and its role in the Internet of Things,” in *Proc. 1st Ed. MCC Workshop Mobile Cloud Comput.*, 2012, pp. 13–16. https://doi.org/10.1145/2342509.2342513

[7] M. Chiang and T. Zhang, “Fog and IoT: An overview of research opportunities,” *IEEE Internet Things J.*, vol. 3, no. 6, pp. 854–864, 2016. https://doi.org/10.1109/JIOT.2016.2584538

[8] Z. Zhou, X. Chen, E. Li, L. Zeng, K. Luo, and J. Zhang, “Edge intelligence: Paving the last mile of artificial intelligence with edge computing,” *Proc. IEEE*, vol. 107, no. 8, pp. 1738–1762, 2019. https://doi.org/10.1109/JPROC.2019.2918951

[9] Y. Mao, C. You, J. Zhang, K. Huang, and K. B. Letaief, “A survey on mobile edge computing: The communication perspective,” *IEEE Commun. Surveys Tuts.*, vol. 19, no. 4, pp. 2322–2358, 2017. https://doi.org/10.1109/COMST.2017.2745201

[10] A. Banks and R. Gupta, “MQTT Version 3.1.1 Plus Errata 01,” OASIS Standard, 2015. https://docs.oasis-open.org/mqtt/mqtt/v3.1.1/errata01/os/mqtt-v3.1.1-errata01-os-complete.html

[11] N. Naik, “Choice of effective messaging protocols for IoT systems: MQTT, CoAP, AMQP and HTTP,” in *Proc. IEEE Int. Syst. Eng. Symp. (ISSE)*, 2017, pp. 1–7. https://doi.org/10.1109/SysEng.2017.8088251

[12] W.-Y. Yi, K.-M. Lo, T. Mak, K.-S. Leung, Y. Leung, and M.-L. Meng, “A survey of wireless sensor network based air pollution monitoring systems,” *Sensors*, vol. 15, no. 12, pp. 31392–31427, 2015. https://doi.org/10.3390/s151229859

[13] M. Mahbub, “Toxic and hazardous gas detection, measurement and monitoring system for safety assurance in home and industrial application of wireless sensor node,” *Eng. Technol. Res.*, vol. 1, no. 3, pp. 89–98, 2019.

[14] American Conference of Governmental Industrial Hygienists (ACGIH), *TLVs and BEIs: Threshold Limit Values for Chemical Substances and Physical Agents*. Cincinnati, OH, USA: ACGIH. https://www.acgih.org/science/tlv-bei-guidelines/

[15] Occupational Safety and Health Administration (OSHA), “Permissible Exposure Limits—Annotated Tables,” U.S. Dept. of Labor. https://www.osha.gov/annotated-pels

[16] Espressif Systems, *ESP32 Series Datasheet*, 2023. https://www.espressif.com/sites/default/files/documentation/esp32_datasheet_en.pdf

[17] P. Warden and D. Situnayake, *TinyML: Machine Learning with TensorFlow Lite on Arduino and Ultra-Low-Power Microcontrollers*. Sebastopol, CA, USA: O’Reilly Media, 2019. https://www.oreilly.com/library/view/tinyml/9781492052036/